MNRAS 000, 1–13 (2017) Preprint 18 April 2017 Compiled using MNRAS LATEX style file v3.0

Separation of pulsar signals from noise using supervised machine learning algorithms

Suryarao Bethapudi[1]⋆ , Shantanu Desai[1]

[1] *Department of Physics, IIT Hyderabad, Kandi, Telangana-502285, India*

arXiv:1704.04659v1 [astro-ph.IM] 15 Apr 2017

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We evaluate the performance of four different machine learning algorithms (ANN, Adaboost, GBC, XGBoost), in the separation of pulsars from radio frequency interference (RFI) and other sources of noise, using a dataset consisting of pulsar candidates obtained from the post-processing of a pulsar search pipeline. This dataset was previously used for cross-validation of the SPINN-based machine learning engine, which was used for the re-processing of the HTRU-S survey. We report a variety of quality metrics from all four of these algorithms. We apply a model-independent information theoretic approach to determine the features with the most predictive power, and also compare with the feature importance results from the machine learning algorithms, wherever possible. We find that the RMS distance between the folded profile and sub-integrations is the most important feature in Adaboost and XGBoost. In the case of GBC, we find that the logarithm of the ratio of barycentric period and dispersion measure to be the most important feature. The information theoretic approach to feature importance yields a ranking very well matched to that based on GBC. For all the aforementioned machine learning techniques, we report a recall of 100% with false positive rates of 0.15%, 0.077%, 0.1%, 0.08% for ANN, Adaboost, GBC, and XGBoost respectively. Amongst all four of these algorithms, we find that Adaboost has the minimum overlap between the error rates as a function of threshold for detection of pulsars and RFI, and based on this criterion can be considered to be the best.

Key words: methods: data analysis stars: neutron

1 INTRODUCTION

Ever since the accidental discovery of 'The Little Green Man' (Hewish et al. 1968), huge amount of manpower and computing resources have been invested in exploring the observable Universe to detect radio pulsars. Pulsars are highly magnetized rotating neutron stars, with misaligned magnetic and rotation axis, which emit pulsed radio emission. Their radio signal is observable when the axis of their emission cone is directed along the line of sight to the observer. Subsequently, they have also been observed throughout the electromagnetic spectrum. An updated review of the various observational properties of pulsars and other kinds of neutron stars can be found in Kaspi & Kramer (2016).

Pulsars have provided a remarkable laboratory for tests and applications of nearly all branches of physics, from condensed matter physics to quantum chromodynamics, and also wide variety of topics in astrophysics spanning stellar evolution, interstellar medium, cosmology etc. (Blandford 1992; Ransom 2013). They can also be used to provide insights on the nature and distribution of dark matter (Baghram et al. 2011; Desai & Kahya 2016). Pulsar observations paved the way for the first discovery of extrasolar planets (Wolszczan & Frail 1992), provided the first indirect evidence for gravitational waves (Taylor 1994), and could also provide evidence for the direct detection of gravitational waves in the nHz regime (Detweiler 1979). Therefore, it is imperative to discover new pulsars in order to harness their tremendous physics potential.

As of April 2017, there are 2613 known pulsars in the ATNF catalog (Manchester et al. 2005). However, the total pulsar population in our galaxy has been estimated to be from $20,000$ to 10^5 (Yusifov & Küçük 2004; Faucher-Giguère & Kaspi 2006; Johnston & Karastergiou 2017). So more than 90% of the total pulsar population still remains to be discovered.

Pulsar searches from modern radio surveys involve sifting through candidates detected from pulsar search pipelines, consisting of either periodicity or single-pulse searches. These pulsar search algorithms are computationally very expensive (although improvements continue to be made to their speed and sensitivity eg. Smith 2016; Cameron et al. 2017). The output of these pipelines yield millions of candidates, out of which a small fraction consists of pulsars and most of the rest are from radio frequency interference (RFI) or other sources of noise (Keith et al. 2010). Many of

⋆ ep14btech11008@iith.ac.in

© 2017 The Authors

these candidates are hand-scanned and manually vetted by domain experts. For current generation pulsar surveys, it has been estimated that about 250 person hours would be required to visually vet the million or so candidates (Lee et al. 2013). Such a manual visual classification of the pulsar candidates cannot be sustained during the SKA era, where we expect to discover about 20,000 new pulsars (Kramer & Stappers 2015). Even though the sifting of real pulsar signals from noise can be facilitated with graphical utilities such as `JREAPER` (Keith et al. 2009), these have limitations and one can easily make mistakes (Bates et al. 2012). Therefore, in order to maximize detection of pulsars in the SKA era, the computational costs in all steps of the pulsar search pipeline should be reduced and human intervention should be minimized at very step. An important step in this process would be to automate the validation of pulsar candidates obtained from pulsar search pipelines as much as possible.

Therefore, for the autonomous removal of true signals from noise, the radio pulsar community has resorted to machine learning to solve this problem (Eatough et al. 2010; Bates et al. 2012; Morello et al. 2014; Zhu et al. 2014; Devine et al. 2016). In most of these papers, the machine learning algorithm used is some variant of Artificial Neural Networks (ANN). Supervised machine learning using ANN was first used in the pulsar community by Eatough et al. (2010) to process 16 million pulsar candidates obtained by reprocessing of data from the Parkes multi-beam survey. Bates et al. (2012) also used ANN in the data-processing pipeline for the High Time Resolution Universe (HTRU) mid-latitude survey. They were able to reject 99% of the noise candidates and detect 85% of the pulsars through a blind analysis. Zhu et al. (2014) used a combination of three different supervised algorithms, namely ANN, Support Vector Machine (SVM), and Convolution Neural Nets in their image recognition based post-processing pipeline dubbed PICS (Pulsar Image-based Classification System) AI. PICS was trained with data from the Pulsar Arecibo L-band Feed Array (PALFA) survey and validated with data from the Green Bank North Celestial Cap survey. From the validation set, PALFA was able to rank 100% of the pulsars in the top 1% of all candidates, while 80% were ranked higher than any noise or interference events. Morello et al. (2014) used neural networks in a pulsar ranking pipeline dubbed Straightforward Pulsar Identification using Neural Networks (SPINN). SPINN used ANN trained on about 1200 observations from the HTRU all-sky survey combined with about 90,000 noise candidates. This pipeline was able to identify all pulsars in the HTRU-S survey with a false positive rate of 0.64% and also helped reduce the amount of candidates to scan by up to four orders of magnitude. In this work, we apply multiple machine learning algorithms on the same dataset, as the candidates used for the cross-validation of the SPINN pipeline were made publicly available [1]. Most recently, Devine et al. (2016) used six different machine learning algorithms to classify dispersed pulsar groups in the second stage of their single-pulse search pipeline, which is also sensitive to other "siblings" of radio pulsars, such as RRATs and FRBs. These six algorithms included ANN, SVM, direct rule learner, standard tree rule learner, hybrid rule-and-tree learner, and Random Forests. The benchmark dataset consisted of over 300 pulsar signals and about 9,600 noise candidates using observations from the Green Bank Telescope. They found that multi-class ensemble tree learners were the most efficient. Therefore, in most of the above papers, some variant of ANN has been used in the post-processing of candidates from the pulsar search pipelines. Here, we start out with ANN and showcase other algorithms which perform better than ANN.

The outline of this paper is as follows. In Section 2, we provide an abridged pedagogic introduction to machine learning, including the definition of various quality metrics used to grade our machine learning algorithms. The pulsar dataset from the HTRU-S survey along with the features used for training is described in Section 3. We discuss the procedure of feature selection prior to running our machine learning algorithm in Section 4. In Section 5, we list the various machine learning algorithms applied to this dataset. Our implementation and results on ANN, Adaboost, Gradient Boosting Classifier, and XGBoost are described in Sections 6, 7, 8, and 9 respectively. We then conclude in Section 10.

2 SUPERVISED MACHINE LEARNING RUDIMENTS

There are two types of supervised machine learning problems, namely *classification* and *regression*. Classification, using machine learning involves determining a relation between the input variables, known as *features* and the output variables, called *class labels*, which is a choice from a pre-defined list of possibilities. A special case of classification involving only two classes is called *binary classification*. The separation of real pulsar signals from noise is an example of *binary classification*. The relationship between the features and class variables is determined using a training set, which is a subset of the data for which we know both the features and the class labels. The training set is usually split into a cross-validation set, used to tune the parameters of the learning algorithm, until we get an optimized performance, as well as a test set used to evaluate the performance of the algorithm with the best tuning parameters obtained from the cross-validation set. The performance on the test set is used as the criterion to select the best algorithm. Supervised machine learning techniques for binary classification are used throughout electromagnetic astronomy at all wavelengths (Ball & Brunner 2010; Cavuoti 2013) and also in neutrino and gravitational wave astronomy (Zevin et al. 2017). We now discuss various metrics to assess the efficacy of various machine learning algorithms.

2.1 Evaluation metrics

In order to grade the performance of various machine learning algorithms, we calculate various performance metrics for each algorithm, which we define in this section. We have used six different metrics to quantify the model performances, viz. Recall, Precision, Accuracy, f_1 score, log loss, and Matthews Correlation Coefficient (MCC).

In addition to these six metrics, we also evaluate the confusion matrix for every algorithm. A confusion matrix is a $c \times c$ matrix, where c is the number of class labels. Each C_{jk} element of the confusion matrix is the number of observations LABELED by the classifier as j, but which belongs to class k. In the special case of binary classification problem, each element of the confusion matrix corresponds to TP, TN, FP, FN, which are defined below (see Table 1). In a multi-class classification setting, such a correspondence isn't there.

(i) **TP**: True Positive. Those observations belonging to the TRUE class, which are correctly labeled by the classifier as POSITIVE.

[1] http://astronomy.swin.edu.au/~vmorello/

	POSITIVE	NEGATIVE
TRUE	TP	FN
FALSE	FP	TN

Table 1. Confusion matrix for a binary classification problem. The rows correspond to the GROUND TRUTH labels and the columns correspond to the labels as classified by the classifier. Definitions of the symbols used here can be found in Section 2.1

(ii) **TN**: True Negative. Those observations belonging to the FALSE class, which are correctly labeled by the classifier as NEGATIVE.
(iii) **FP**: False Positive. Those observations belonging to the FALSE class, which are incorrectly labeled by the classifier as POSITIVE. Hence, the name False Positive. This type of error is also called Type-I error.
(iv) **FN**: False Negative. Those observations belonging to the TRUE class, which are incorrectly labeled by the classifier as NEGATIVE. This is also referred to as Type-II error.

Now, we define each of the six metrics for all the algorithms used in this paper. A tabular summary of all of these metrics can be found in Table 2.

Recall (also known as Sensitivity) is defined as

$$\frac{TP}{TP+FN}$$

A Recall score of one implies that all positively classified observations belong to the TRUE class. This metric is given prominence, when misclassifying a TRUE class observation has a higher cost than misclassifying a FALSE class observation, or in other words when Type-II error has a higher penalty than Type-I error. This is the exactly the case we need to consider for the pulsar separation problem discussed in this paper.

Precision is defined as

$$\frac{TP}{TP+FP}$$

Precision is the ratio of TRUE class observations labeled correctly to the total number of observations, which were labeled as positive (irrespective of whether they are true positives or not).

Accuracy is considered the most important metric in any classification problem and is defined as

$$\frac{TP+TN}{TP+TN+FP+FN}$$

f_1 **score** also known as F-measure or balanced F-score (Bishop 2006) is defined as

$$\frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

In other words f_1 score is twice the harmonic mean of Precision and Recall. It provides an alternate measure of accuracy.

log loss also known as logistic loss or cross-entropy loss is defined as

$$\text{logloss} = -\frac{1}{\mathrm{N}}\sum_{\mathrm{i}=1}^{\mathrm{N}}(\mathrm{y_i}\log(\mathrm{p_i}) + (1-\mathrm{y_i})\log(1-\mathrm{p_i}))$$

where y_i are class labels and p_i are corresponding probabilities. Interested readers are advised to refer to Bishop (2006) for more details.

MCC is an acronym for Matthews correlation coefficient (Matthews 1975). It is a measure of the quality of binary classification. It quantifies the correlation between observed and predicted binary classifications. MCC is defined as

$$\frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$$

Additional details about MCC can be found in Baldi et al. (2000).

3 PULSAR DATASET

3.1 Pulsar Detection

We briefly outline the steps involved in the detection of pulsars from raw radio data using periodicity searches. More details on each of these steps can be found in Lorimer & Kramer (2004) and references therein. The first step in the pulsar search pipeline involves de-dispersion, which corrects for the dispersion of the interstellar medium. The amount of dispersion is governed by the integrated electron density along the line of sight, also known as the dispersion measure (DM). A brute force search algorithm for pulsars involves a grid search in DM and acceleration space. For each putative DM and acceleration, periodic signals are searched for using FFT techniques and potential pulsar candidates are then selected from these searches. These candidates are then manually inspected to veto the noise sources, to create the final list containing only pulsars.

3.2 HTRU Survey

The dataset we have used in this manuscript for testing various algorithms is obtained through the reprocessing of data from the Southern High Time Resolution mid-latitude survey (Keith et al. 2010) by the `PEASOUP` pipeline. Briefly, the Southern HTRU Latitude survey (HTRU-S) is a survey of the southern sky ($30^\circ > l > -120^\circ; |b| < 15^\circ$), carried out using the 13-beam receiver on the Parkes Radio Telescope to detect radio pulsars and also short-duration radio transients. A similar survey of the Northern skies has also been conducted using the Effelsberg radio telescope with a seven-beam receiver (Ng 2013). More details of the HTRU-S survey strategy and specifications can be found in Keith et al. (2010). This data has been reprocessed using a GPU-based `PEASOUP` pipeline with acceleration searches turned on (Morello et al. 2014). The pipeline searched for dispersion measures between 0 and 400 cm^{-3} pc and accelerations between -50 and 50 $\mathrm{m/s^2}$. This search returned about 4.3 million candidates. Among these, 1200 were real pulsars and to these real pulsars, 90,000 non-pulsar observations (from RFI and other sources of noise) were added in order to test a machine-learning based solution (called SPINN) to the problem of separating pulsar signals from noise (Morello et al. 2014). The pulsar candidates encompass a diverse range of periods, duty cycles, and signal to noise ratios. It is unlikely that there are undiscovered pulsars in the set of noise candidates. More details of the pre-processing steps can be found in Morello et al. (2014). This dataset (with the pulsar and noise candidates tagged for each event), was made publicly available by Morello et al. (2014) for the development and testing of various machine learning algorithms by the wider scientific community. Our goal is to find a machine learning solution with the maximum efficiency in detecting pulsars and minimum contamination from noise candidates. In machine learning parlance as discussed earlier, we are tackling a binary classification problem.

Table 2. Definitions of various metrics used in this letter. The 'Range' column corresponds to the range of values each metric can take. The best value a metric can take is **boldfaced**.

Metric	Formula	Description	Range	
Recall	$\frac{TP}{TP+FN}$	Ability to "remember" all the positive samples.	**1**	0
Precision	$\frac{TP}{TP+FP}$	Ability to "detect" negative samples.	**1**	0
Accuracy	Refer Sub. Section 2.1	Ability to make no error of any type in classification	**1**	0
f_1 **score**	$\frac{2*\text{Precision}*\text{Recall}}{\text{Precision}+\text{Recall}}$	Harmonic mean of Precision and Recall.	**1**	0
log loss	Refer Sub. Section 2.1	Cross entropy loss. Closely related to KL divergence.	**0**	∞
MCC	Refer Sub. Section 2.1	Measure of correlation between label and features.	**1**	-1

Each Pulsar candidate is described in a PulsarHunter Candidate XML file, having the extension `.phcx`. See Figure 3.2 to see the schema of the xml tree.

For all the machine learning algorithms herein, we have used the same feature set as defined by Morello et al. (2014). They are listed below, and henceforth we shall denote each feature with a corresponding numerical index (called *feature index*, f_{ID}), using 0-based indexing.

0. SNR of the folded profile
1. Intrinsic equivalent duty cycle of the pulse profile
2. Ratio between barycentric period and dispersion measure *(log scale)*
3. Validity of dispersion measure
4. Persistence of signal in the time-domain
5. RMS distance between the folded profile and the sub-integrations

We defer details of the above features to Morello et al. (2014). All of these were scaled to have zero mean and unit variance. There are only 1196 pulsar observations in the dataset, whereas there are $100,000$ RFI observations. To account for this class imbalance, we emulate the procedure followed by Morello et al. (2014) and replicate the pulsar-only candidate set sufficient number of times, so as to have the ratio of true positives to true negatives equal to 1:4.

4 FEATURE SELECTION

While carrying out an in-depth feature exploration is beyond the scope of the paper, we briefly delve into the paramount task of feature selection, prior to running our machine learning algorithms. In any supervised algorithm, the features that the model feeds on largely determine the accuracy of the model. An information-theoretic way to quantify the usefulness of a feature is to compute the Mutual Information between a feature and class label (Macedo et al. 2017).

However, since this mutual information doesn't include the correlation between the features, we account for it by calculating the difference between the two. This is the crux of the *Max. Relevance and Min. Redundancy* based Feature Selection or mRMR approach (Peng et al. 2005; He et al. 2013). Here, we have used the mRMR algorithm available online[2] to compute MRMR based feature selection. Since, the code initially converts the 'contiguous' type variable to 'categorical' type variable by using two thresholds of the form $\mu \pm k * \sigma$ where μ, σ are the mean and standard deviation of the feature observations to be categorized, and k is an integer, which allows us to choose the thresholds about the mean.

Since, k is an ad-hoc parameter of choice, which can be arbitrarily tuned, we have computed the mRMR based ranking of features in Table 3 for different values of k. An interesting point to notice here is that for the most likely choices of k, i.e for $k = 2$ and $k = 3$, our ranking of features is closely similar to the ranking derived from Gradient Boosting Classifier. See Section 6 later in this manuscript, for more details about this point.

[2] http://home.penglab.com/proj/mRMR/

Table 3. mRMR based feature ranking. See Section 4 for more details about k. Entries in the table (columns 2-6) are f_{ID} as defined in Section 3.2. For most likely choices of k, we find that our ranking is similar to that obtained from Feature Importance in the case of Gradient Boosting Classifier (Section 8). It is also evident that based on MRMR feature, $\#4$ (the persistence of signal in the time domain signal) is the best feature.

Rank	$k=0$	$k=1$	$k=2$	$k=3$	$k=4$
1^{st}	0	0	4	4	0
2^{nd}	1	1	2	2	2
3^{rd}	2	2	1	5	1
4^{th}	3	5	5	1	4
5^{th}	5	3	0	0	5
6^{th}	4	4	3	3	3

We note that MRMR is not a machine learning algorithm or related to any particular one, but is a statistical diagnostic tool to rank the features based on the Mutual Information and correlations. It can be run in a stand-alone mode and is completely decoupled from any machine learning algorithm. The decisive power of MRMR becomes more evident, when there is a plethora of diverse features, out of which we have to select 'n' most powerful features. MRMR is one of the many pre-processing tools, which can be used to select a subset of important features on which any machine learning algorithm can be trained on. While there are machine learning algorithms, which can independently compute feature importances, ranking obtained by these are not uniform and are algorithm-specific. As we shall see in later on in the feature importances plots in Figures 4 and 6, ranks given by Adaboost (Section 7) or XGBoost (Section 9) and GBC (discussed in Section 8) do not match. Therefore, we have computed the MRMR based ranking and later show that the ranking obtained in MRMR agrees closely with the same from GBC.

5 MACHINE LEARNING METHODS

In this work, we have considered four different algorithms to this pulsar classification problem :

(i) Artificial Neural Network (ANN)
(ii) AdaBoost
(iii) Gradient Boosting Classifier (GBC)
(iv) eXtreme Gradient Boosting (XGBoost)

To the best of our knowledge, among these four algorithms, only ANN has been used before in the core pulsar search pipelines or during post-processing of pulsar candidates. Although, we are only reporting the aforementioned techniques, we also tried the

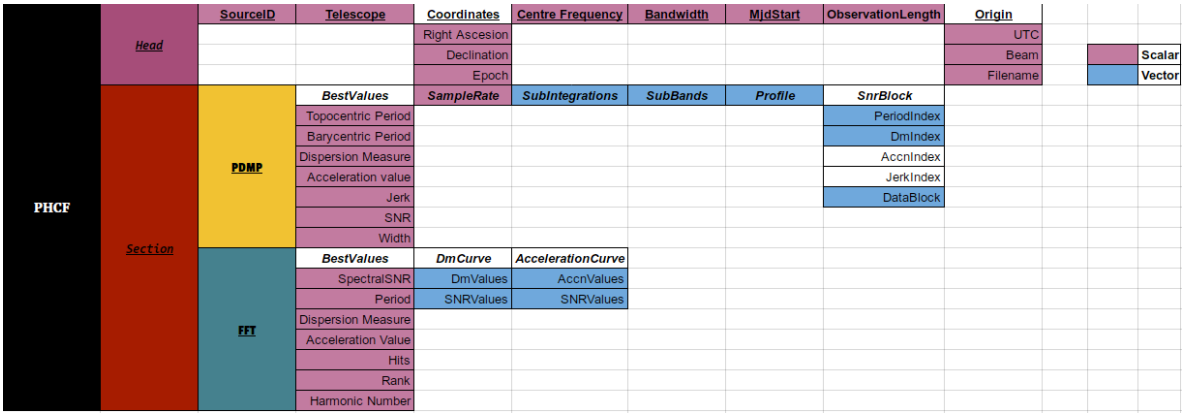

Figure 1. Schema of the XML tree containing details of the pulsar dataset from the HTRU survey, described in detail in Morello et al. (2014) and made publicly available. Each candidate file has the above structure.

Order	**0**	**1**	**2**	**3**	**4**	**5**
0	1.00	-0.34	0.68	-0.02	**0.86**	*-0.46*
1	-0.34	1.00	*-0.66*	**0.07**	-0.33	-0.04
2	0.68	*-0.66*	1.00	-0.32	**0.76**	-0.35
3	-0.02	0.07	*-0.32*	1.00	-0.07	**0.08**
4	**0.86**	-0.33	0.76	-0.07	1.00	*-0.55*
5	-0.46	-0.04	-0.35	**0.08**	*-0.55*	1.00

Table 4. Correlation matrix of our dataset where i^{th}, j^{th} term correspond to the Pearson Correlation coefficient (Press et al. 1992) of f_i, f_j features. We have also performed normalization to have variance $= 1$. 'f_{ID}' refers to Feature Index given in Section 3.2. Numbers appearing in **boldface** and *italicized* indicate the row-wise maximum and minimum respectively.

Table 5. This table shows the number of observations used for training and testing in all the folds. Please note that this distribution is same for all machine learning algorithms discussed here. `#Pulsar` denotes the number of pulsar observations and `#RFI` denotes the number of RFI observations.

Used for	Training		Test	
Fold	`#Pulsar`	`#RFI`	`#Pulsar`	`#RFI`
1^{st} fold	15098	58453	7550	29227
2^{nd} fold	15099	58453	7549	29227
3^{rd} fold	15099	58454	7549	29226

Naive Bayes algorithm. However, the performance of Naive Bayes was unsatisfactory. One of the core premise of the Naive Bayes algorithm is conditional independence between the features for a given target label. We find that because of the significant amount of correlation between the features (cf. Table 4), Naive Bayes fails on this dataset (Zhang 2004).

Therefore, we do not report the results from the Naive Bayes algorithm. We now describe in detail our implementation of the four algorithms and present results from the same. All the four algorithms have been implemented in Python using the `scikit-learn` module (Pedregosa et al. 2011). For each of the above algorithms, we have performed stratified three-fold cross-validation. More details on how the data is split into a training and validation sample, including the number of signal and noise events in each fold of cross-validation are provided in Table 5.

6 ANN

Artificial Neural Network (ANN) is a popular tool in the machine learning community. Complete details of ANN are beyond the scope of this paper and can be found in Bishop (2006) and Hastie et al. (2001). Most machine learning algorithms used by the pulsar community utilize some variant of ANN. ANN is also widely used in optical astronomy (Sadeh et al. 2016), in particular for photometric redshift estimation (eg. Desai et al. 2012). The hyper-parameters used for training the ANN are specified in Table 6. As we can see from Table 6, learning rate is the most important hyper-parameter, where learning rate is defined as the hyper-parameter which quantifies the magnitude of distance we want to go in the direction of its gradient.

We have used the `SKNN` Python module to implement this Neural Network. We have performed three-fold stratified cross-validation, initializing the Neural Network with *Random Seed* equal to 1234. The loss function (objective function) to be minimized (maximized) in ANN is always a function of very high dimensionality. Furthermore, this loss function is non-convex. Thus, we cannot have a closed form solution to the minimization problem and we must resort to gradient-based solutions.

In all such gradient-based solutions, determining the optimal learning rate is often the trickiest problem. If we choose too high a value of this hyper-parameter, we might never reach local minima (maxima) but just pass right by it, whereas too low a value of the learning rate might increase the convergence time. At times, finding an optimal parameter becomes so difficult that switching to alternate machine learning solutions may seem more prudent.

In order to mitigate this issue, we have used AdaGrad (ADAptive GRADient) as our gradient descent optimization algorithm (learning rule), instead of the default Stochastic Gradient Descent (SGD) in `SKNN`. In the SGD optimization algorithm, the learning rate is either constant or is reduced after a fixed number of epochs. AdaGrad instead uses a different learning rate for each parameter of our model and regulates the learning rate on the fly. It adjusts the base learning rate ('original learning rate') by dividing by the square-root of the running sum of squares of the gradients, corresponding to the parameter for which the base learning rate has to be adjusted. There is an additional fudge factor added to this square-root to ensure that the division operation does not throw an exception. More details about AdaGrad can be found in Duchi et al. (2011).

Morello et al. (2014) employ a single hidden layer with eight

Table 6. Hyper-parameters used in ANN training of the pulsar dataset. More details about these parameters can be found in any application of ANN (for eg. Hastie et al. 2001; Bishop 2006; Srivastava et al. 2014; Duchi et al. 2011).

Property	Value
Batch Size	500
Number of hidden layers	1
Weight decay	10^{-5}
Dropout	0.4
Number of hidden nodes	10
Activation function	RLeU
Output Layer	Softmax
Learning Rule	AdaGrad
Learning Rate	10^{-2}

hidden units of Neural Network with $\tanh$ as the output layer activation function. They have trained using 'mini-batch' approach, wherein the weight updates happen after a batch (subset) of the data has been passed. While the architecture of ANN we have implemented is the same as the one implemented by Morello et al. (2014), with the exception of two additional hidden neurons in our hidden layer, we have used a different learning rule (AdaGrad), activation function (RLeU), and output layer of softmax. Moreover, we also used dropout (Srivastava et al. 2014) as a measure against over-fitting.

In addition, they record the 'score' (output at the output layer) and plot the error rates as a function of the score threshold. We have made a similar plot in Figure 9. Furthermore, we also plot the confusion matrix in Figure 2.

We have trained an ANN and achieved 100% Recall. We report an Accuracy of 0.984 with Precision of 0.93, f_1 score of 96.3%, log loss of 53.1%, and MCC of 95.5%. We also report FPRs of $0.15\%, 0.12\%, 0.05\%$ for values of P_T equal to $0.5, 0.7, 0.9$ respectively, where P_T corresponds to the threshold probability. We label observations with $Prob\{pulsar\} > P_T$ as pulsars. The error rates (for both pulsars and noise candidates) as a function of P_T for ANN can be found in Figure 9. The corresponding error rate plots for other models can be found in Figures 10, 11, and 12. See Table 9 for a summary of all the quality metrics we calculate for ANN and in-depth comparative assessment with other algorithms to be described next. The Precision, Recall, FPR for selected threshold probabilities for all ANN (and also others) can be found in Table 10.

ANNs are termed as Universal approximators; they can be trained to model any function given sufficient data. In order to capture all of the intricate non-linearities of the data, one can cascade any number of these hidden layers. In principle, we can achieve any amount of accuracy just by increasing the number of layers but this makes the training process computationally intensive, potentially leading to over-fitting as the number of weights increases exponentially but the data remains the same. At this point, it is wise to invoke Occam's Razor, which points us in the direction of the simplest model.

7 ADABOOST

AdaBoost, an acronym for Adaptive Boosting, takes an ensemble of 'weak learners' and repeatedly trains them on copies of the original dataset, all the while focusing on the 'difficult' (or outlier) data points *(hence the word 'Adaptive')* (Freund & Schapire 1995). This, in turn, ensures that our ensemble *(boosted classifier)* performs drastically better on the test data. A natural choice of weak learners are decision trees having the depth of one level. *Such trees are called 'decision stumps'*. Refer Section 2 of Mayr et al. (2014) for an overview of AdaBoost. Adaboost has been shown to outperform other machine learning algorithms for many different supervised classification problems in optical astronomy (Hoyle et al. 2015; Sevilla-Noarbe & Etayo-Sotos 2015; Elorrieta et al. 2016; Acquaviva 2016; Zitlau et al. 2016).

Table 7. Hyper-parameters used in AdaBoost training.

Property	Value
Number of trees	500
Max. No of Leaf Nodes	3

We use the same implementation of Adaboost as in Hoyle et al. (2015), where it was applied to photometric redshift estimation of SDSS DR10 data and shown to be more robust against outliers and provided better match with spectroscopic redshifts than ANNs. We have used 9876 as the random seed and the hyper-parameters used are specified in Table 7. We have defined our AdaBoost model using two properties (specified in Table 7), namely, the number of trees and the maximum number of leaf nodes. The former corresponds to the number of estimators and the latter to the base estimator's hyper-parameter. Please note that Adaboost is a meta classifier which takes `N copies` of the base estimator (weak learner) and trains them on the same feature set but with different weights assigned. These `N copies` correspond to the number of estimators. Moreover, the maximum number of leaf nodes restrict the maximum height of the decision tree, which is the base classifier in this case.

See Belanich & Ortiz (2012), which extols the power of Adaboost on mathematical grounds. Given this, when we perform three-fold stratified cross-validation, we get extremely good results with 100% recall during all the three iterations with very low FPR. The simplest explanation as to why AdaBoost works so well would be that it focuses more on data points, which are incorrectly classified. It assigns more weight to those data points, that are incorrectly classified during the training process *(hence the name adaptive)*. This, in turn, ensures that even the tricky data points are used during training.

Another advantage of AdaBoost is that we can calculate the 'relative feature importance' for each feature. This plays an important role in feature selection. Relative feature importance is measured by how much error that feature reduces each time it was used as a criterion in the weak classifier ('decision stump'). In other words, if I use this feature, how well I am able to segregate signals from noise. 'Wellness' is measured by the difference in error before and after splitting. Larger the change, more 'powerful' is this feature. See Hoyle et al. (2015) and Chapter 10 of Hastie et al. (2001) for a deeper understanding of feature importance calculation in Adaboost. Our results for feature importance in Adaboost can be found in Table 4 We notice that relative feature ordering is more or less the same in all the three iterations. This, in turn, implies that Adaboost always chooses the same order of features to achieve it's superior performance (see Table 9). One also notices that feature #5, has a higher feature importance by a good margin compared to the other features. Hence, we find that feature #5 (RMS distance between the folded profiles and sub-integrations) as the best feature for Adaboost. (As we shall see later, in the case of GBC, we find that feature #5 isn't as important as perceived in Adaboost.)

For 100% Recall, Adaboost achieves a Precision of 99.3%, False positive Rate of 0.0077%, Accuracy of 99.8%, f_1 score of

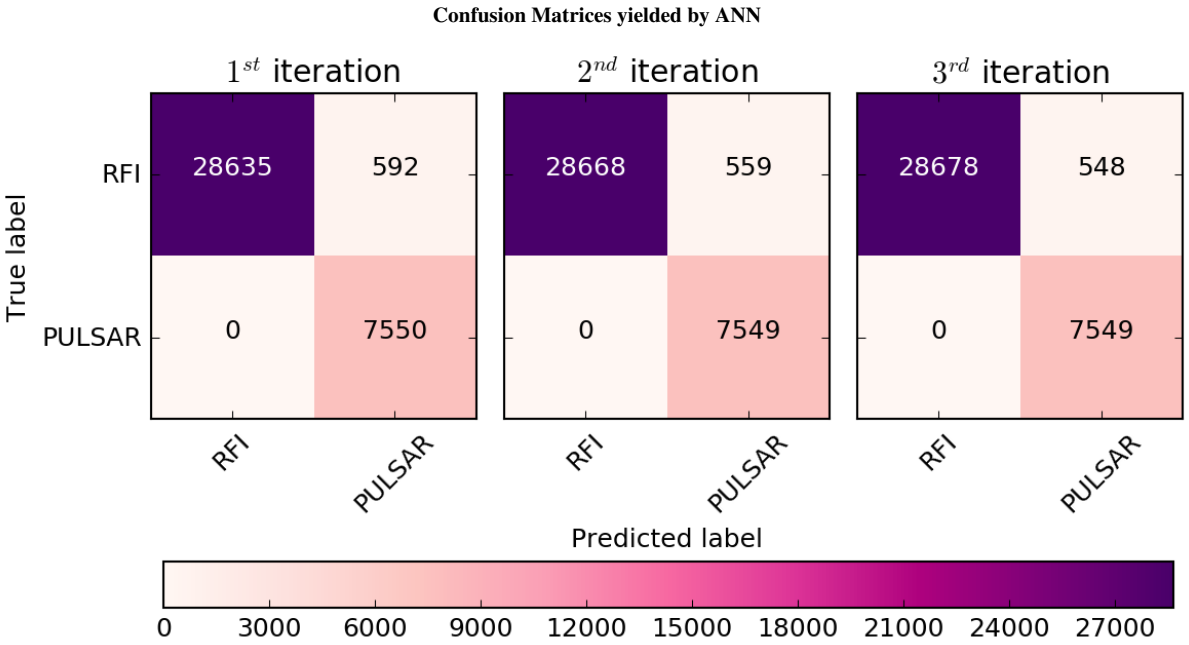

Figure 2. Confusion matrix obtained after running ANN. We have achieved 100% Recall in all the three folds. Details of the hyper-parameters used can be found in Table 6.

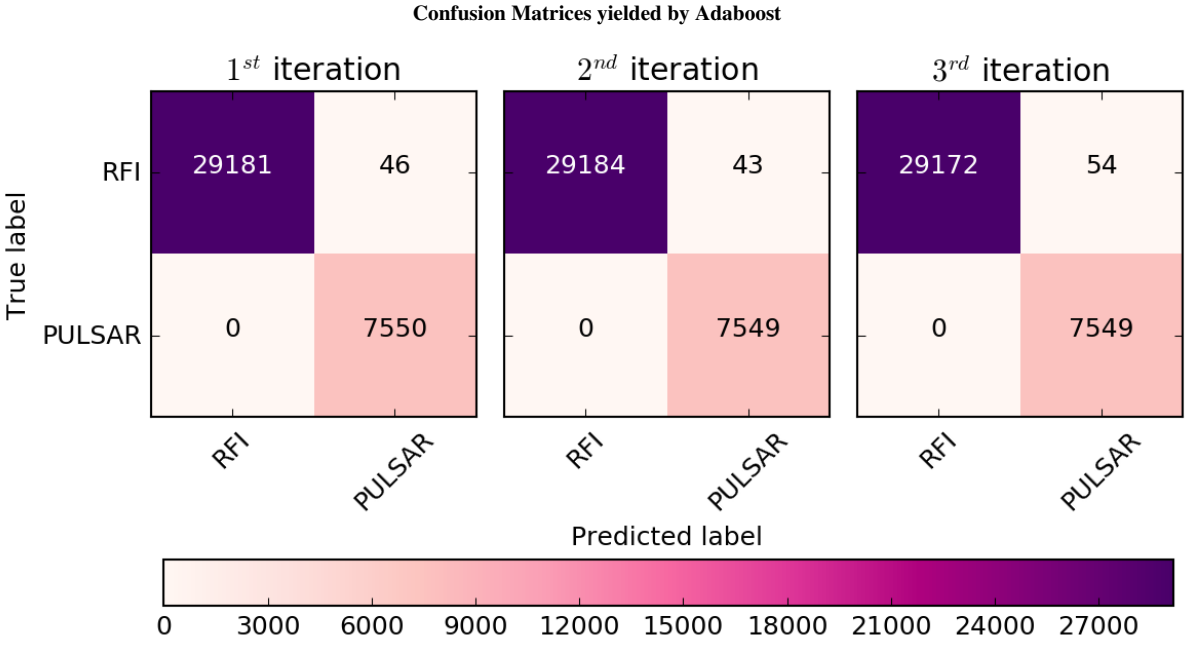

Figure 3. Confusion matrix obtained after running Adaboost. We have achieved 100% Recall in all the three folds. Details of the hyper-parameters used can be found in Table 7.

99.6%, log loss of 4.4%, and MCC of 99.6% (cf. Tables 9 and 10). The error rate as a function of threshold probability can be found in Figure 10. We find FPRs of $0.077\%, 0.013\%, 0.0\%$ for P_Ts of $0.505, 0.515, 0.530$ respectively.

In summary, we find that Adaboost performs better than ANN and GBC (next section) and competes closely with XGBoost (Section 9). It achieves the most desirable scores in every metric we computed. Adaboost, in general, works not only with decision stumps (as done in this paper), but also with any classifier. It is not a model in itself but is an enhancement which maximizes the learning capabilities of base learners.

8 GRADIENT BOOSTING CLASSIFIER

Gradient-based learning is the go-to algorithm, when there is no closed form analytical expression for minimizing the loss function (or maximizing the objective function). Here, we posit our loss

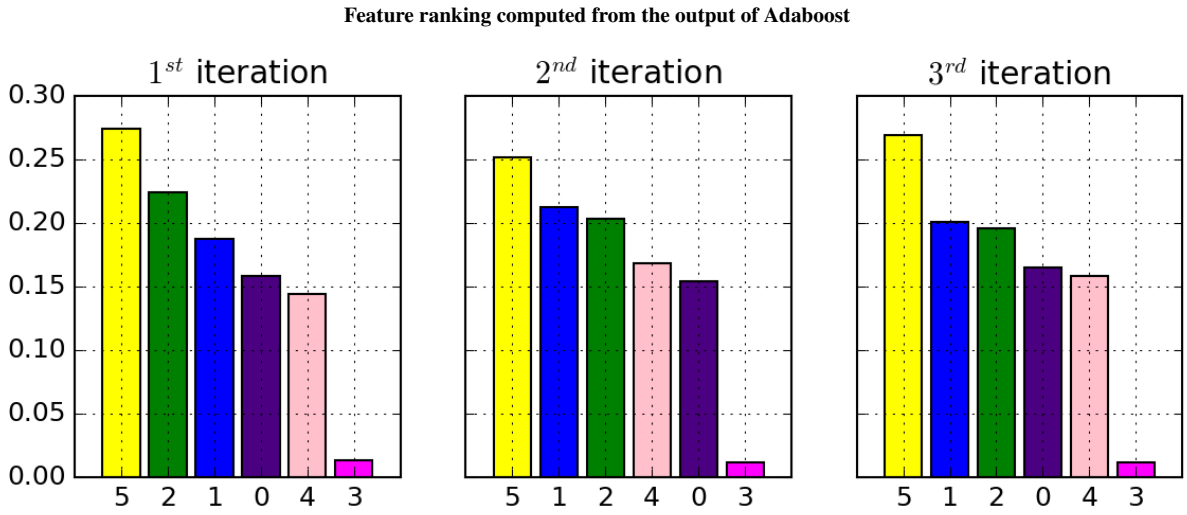

Figure 4. Feature ranking from three-fold cross-validation after running Adaboost. Each feature has a unique color assigned to it. The mapping used here is kept uniform throughout the paper. Larger the value, more is the predictive power of the feature. The most important feature is feature # 5, viz. the RMS distance between the folded profile and sub-integrations. See Figure 6 for the corresponding plot from the output of GBC.

function to be a functional of the 'data' and the 'parameters'. *These parameters are what we are interested in.* Then, we calculate the gradient with respect to each parameter. This is followed by multiplying by the 'learning rate' (the most important parameter in gradient based learning algorithms and same as in ANN), followed by subtracting each parameter (not hyper-parameter) from it's own corresponding product from itself. The "learning rate" parameter in gradient boosting controls how strongly each tree tries to correct the mistakes of the previous trees. A higher learning rate implies that each tree can make stronger corrections, allowing for more complex models. Intuitively, with this procedure we determine those parameters, which minimize our loss function.

Gradient Boosting Classifier (GBC) builds a regression tree (since binary classification) in an 'additive' fashion. *They hail from a broad class of models called Generalized Boosting Models.* Consider an ensemble of weak learners ('trees of same depth'). At each 'stage', you add a weak learner, which is trained on the errors of the model, to the model along with a multiplicative constant. See Friedman (2001); Tramacere et al. (2016) for a thorough understanding of Gradient Boosting Machines. Gradient boosting has also been widely used in astrophysics (Möller et al. 2016; Tramacere et al. 2016) and in particle physics (Yang et al. 2005). More details about Gradient boosting machines can be found in Natekin & Knoll (2013).

We have used the same hyper-parameters in training GBC as in the case of AdaBoost and can be found in Table 7. In our `scikit-learn` implementation of GBC, we have used a learning rate of 0.1.

An interesting thing to note here is that while Adaboost and GBC have the same hyper-parameters, the performance of GBCs is slightly inferior compared to Adaboost. However, please note that they both have a comparable accuracy of about 99.8% (see Table 9).

Similar to AdaBoost, we show the feature importance plots here in Figure 6. Note that feature importance is calculated in the same way as in AdaBoost, i.e. by calculating the improvement in the error we receive after using an additional feature.

For 100% recall, GBC achieves precision of 99.2%, f_1 score of 99.5%, log loss of 5.6%, MCC of 99.4% and accuracy of 99.8% (cf. Table 9). Additionally, GBC obtains FPRs of $0.1\%, 0.09\%, 0.08\%$ for P_Ts of $0.84, 0.9, 0.94$ respectively (cf. Table 10). The error rate as a function of threshold probability can be found in Figure 11. We stress that compared to Adaboost, the performance of GBC is slightly inferior, albeit the accuracies achieved by both are about the same.

9 XGBOOST

XGBoost (Chen & Guestrin 2016) is yet another tree-based model which gained much popularity right since it's inception in the machine learning community in 2016. This model has been the winner in many machine learning challenges and our problem is no exception to that. One of the key features of XGBoost is the scalability it offers. This means that it can be run on distributed systems and also work on very large datasets with ease. While our dataset is not as huge as those on which XGBoost is usually run, our problem serves as a proof-of-principles demonstration of having an end-to-end pulsar separation methodology.

XGBoost, is an acronym for eXtreme Gradient Boosting, and works in the same way as Gradient Boosting with the addition of Adaboost-like feature of assigning weights to each observation. These weights, in turn, help in deciding the tree split. This is a very high level comparison in the sense that Adaboost has an objective/loss function of a specific type. *In classification using Adaboost, most often exponential loss (as in our case) is used, whereas Gradient Boosting and XGBoost have customizable objective/loss functions.* In astrophysics, XGBoost was recently applied for the classification of unknown point sources in the Fermi-LAT catalog (Mirabal et al. 2016) and also for studying the stability of planetary systems using physically motivated features with the help of supervised learning (Tamayo et al. 2016).

As done earlier, we carried out three-fold cross-validation using the hyperparameters tabulated in Table 8. We achieved 100% recall with Precision of 0.993 and Accuracy of 0.998. We also achieved f_1 score of 0.996, log loss of 0.049 and MCC of 0.995 (cf. Table 9). Please note that the performance is closer to Adaboost. We also report FPRs of 0.08%, 0.07% and 0.06% for P_T of 0.85, 0.91

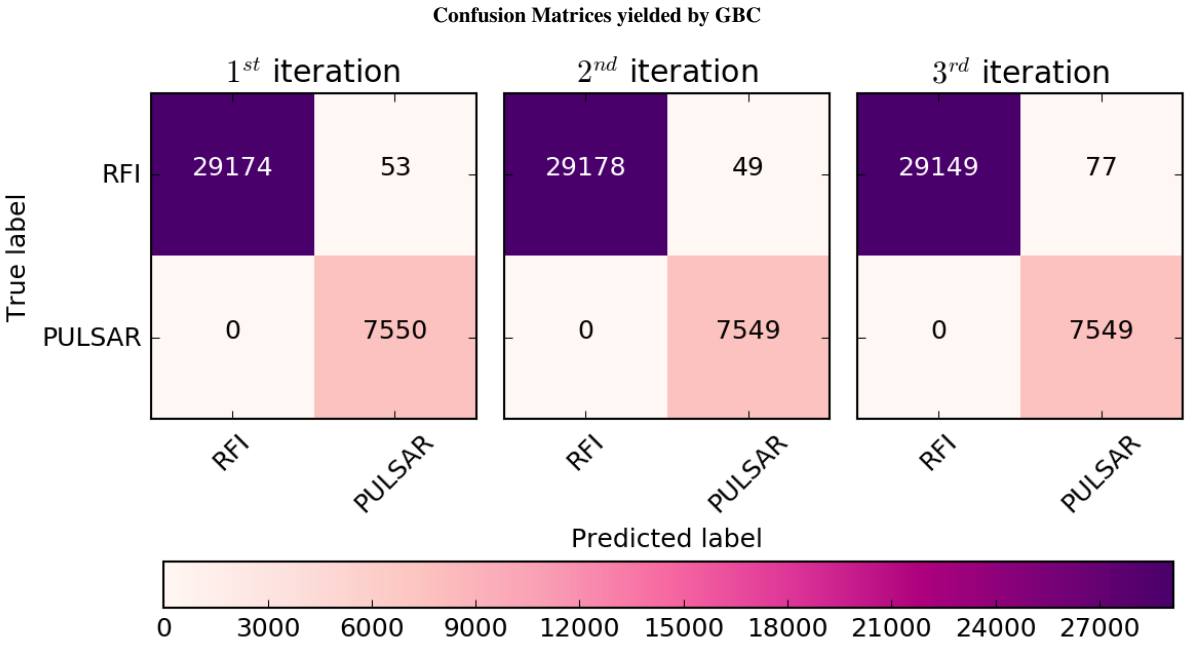

Figure 5. Confusion matrix obtained using Gradient Boosted Classifier. We have achieved 100% Recall in all the three folds. Each feature has a unique color assigned to it. The mapping used here is kept same throughout the paper. The list of hyper-parameters used can be found in Table 7. Note that we have used the same hyper-parameters as in the case of Adaboost.

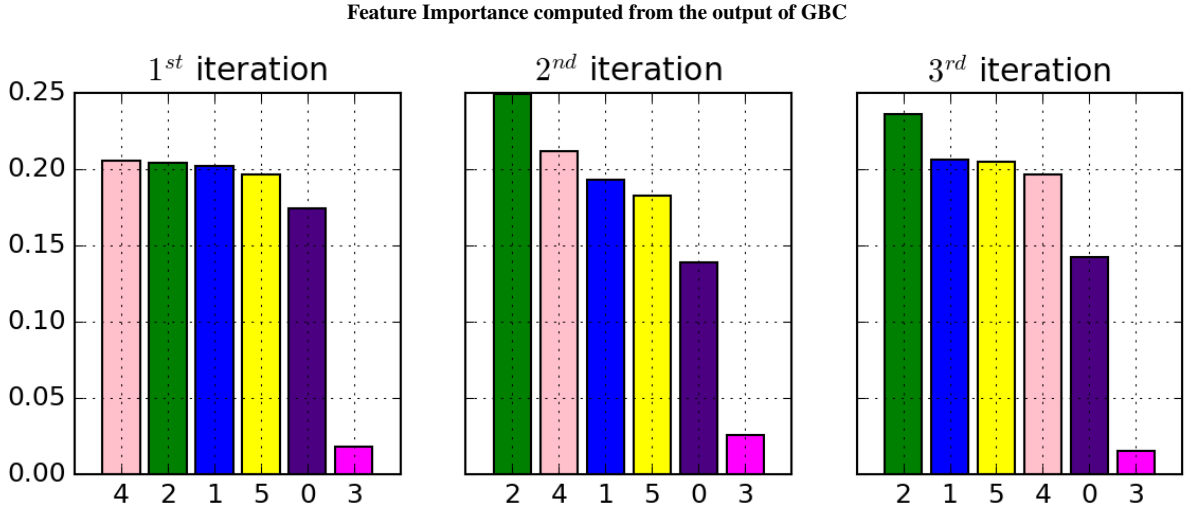

Figure 6. Feature ranking after each iteration of three-fold cross-validation. We have used the same color code for each feature. Larger the value, more powerful is the feature in prediction. The rank of all he features in descending order (in terms of feature importance), in the case of GBC, is the same as we got in mRMR approach (cf. $k = 3$ column of Table 3. Also, notice the difference in feature importance compared to that from Adaboost (Figure 4).

and 0.94 corresponding to recall of 1.0, 0.996, and 0.985 respectively (cf. Table 10). Similar to the other three algorithms we also plot the feature importance, confusion matrix, as well as error rates as a function of threshold probabilities. The confusion matrix can be found in Figure 7. The feature importance plots can be found in Figure 8. As we can see, the relevant rank of all the features is the same as in Adaboost. The error rate as a function of threshold probability can be found in Figure 12. The performance of XGBoost is very close to that of Adaboost despite the former being a Gradient Boosting mechanism unlike GBC. Not only the performance, but also the metrics and feature rankings returned by Adaboost and XGBoost strongly agree. (c.f. Figures 4, 8 and Table 9.). The false positive rates are also comparable to Adaboost (see Table 10).

Table 8. Hyper-parameters used in training XGBoost Classifier.

Property	Value
Number of trees	500
Learning rate	0.1

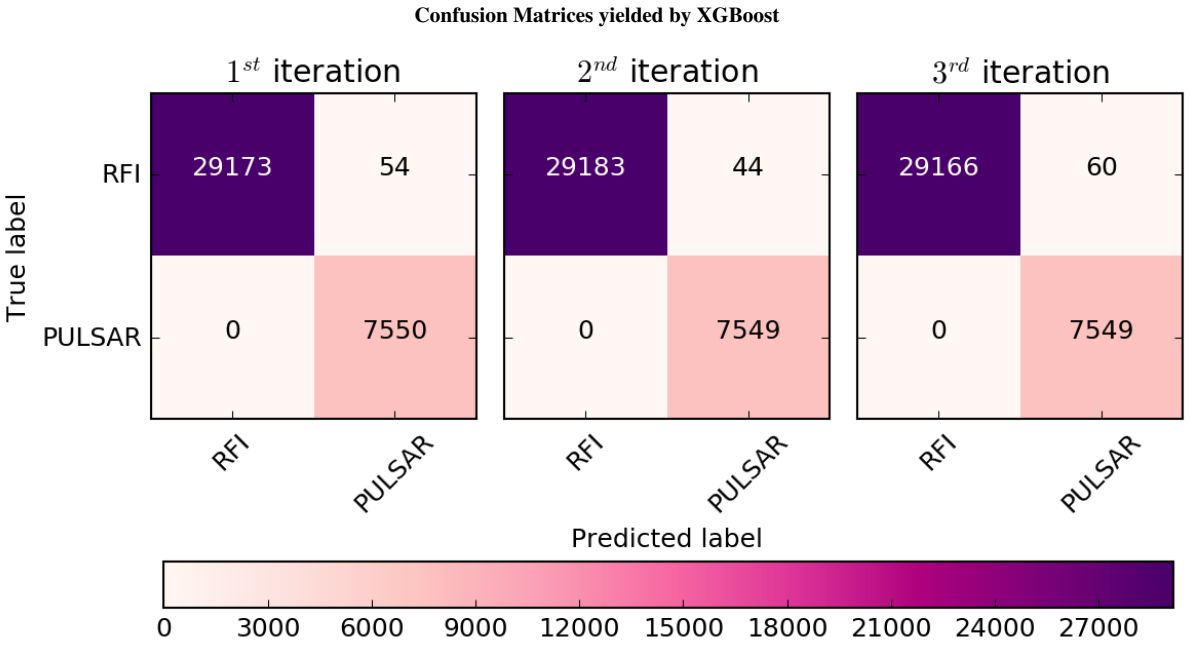

Figure 7. Confusion matrix obtained after running XGBoost. We have achieved 100% Recall in all the three folds. Details of the hyper-parameters used can be found in Table 8.

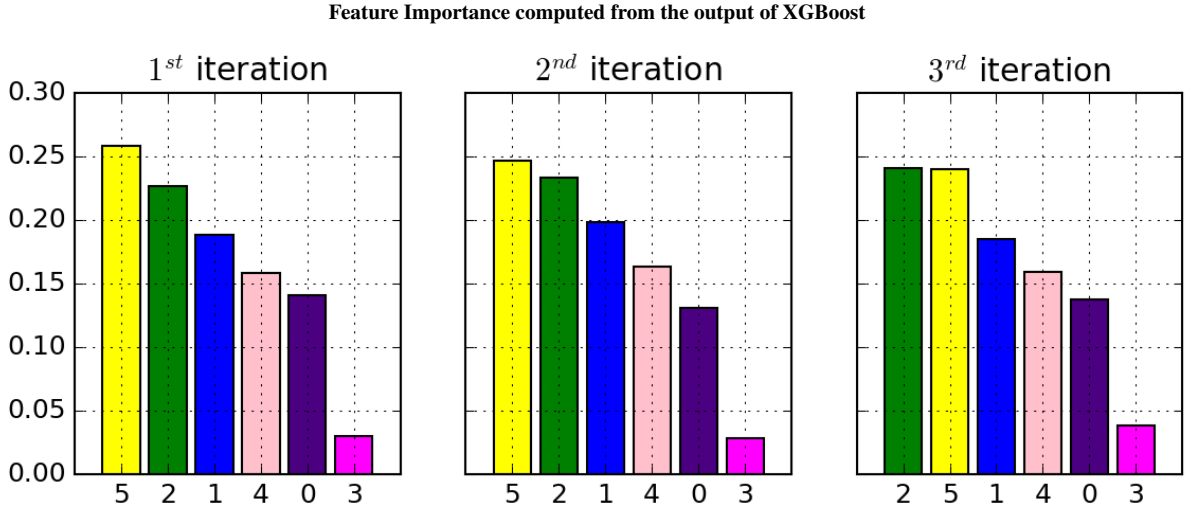

Figure 8. Feature ranking after each iteration of three-fold cross-validation. We have used the same color code for each feature. Larger the value, more powerful is the feature in prediction. The rank of all the features in descending order (in terms of feature importance), in the case of XGBoost, is the same as we got in Adaboost (Figure 4).

Table 9. Metric scores of various models we considered. For definitions of each of these metrics, see Table 2. Column maximums have been **boldfaced**. A point to consider here is that while almost all metric scores of Adaboost and GBC are very close, they seem to differ by a big margin in the case of 'log loss'.

Model	Recall	Precision	Accuracy	f_1 score	log loss	MCC
XGBoost	**1.0**	**0.993**	**0.998**	**0.996**	0.049	0.995
Adaboost	**1.0**	**0.993**	**0.998**	**0.996**	**0.044**	**0.996**
GBC	**1.0**	0.992	**0.998**	0.995	0.056	0.994
ANN	**1.0**	0.930	0.984	0.963	0.531	0.955

10 DISCUSSION AND CONCLUSIONS

We have run multiple supervised machine learning algorithms to evaluate the efficacy and robustness of the separation of true pulsar signals from noise candidates, using data from the HTRU-S survey, a part of which has been made publicly available by Morello et al. (2014). We have reported our implementation and results from four of these algorithms, namely ANN, Adaboost, Gradient Boosting Classifier, and XGBoost, all of which achieved 100% Recall. Before applying these algorithms we also implemented feature selection using the mRMR approach. We evaluated various quality metrics for all four of these algorithms, which can be found in Table 9. The Precision, Recall, and False Positive Rates for all these algorithms can be found in Table 10. Furthermore, we have shown the feature ranking plots in the case of AdaBoost, Gradient Boosting Classifier, and XGBoost. As shown in Figs. 4 and 8, we note that $f_{ID} = 5$, viz. the RMS distance between the folded profile and sub-integrations is the most important feature in Adaboost and XGBoost. In the case of GBC, we note that $f_{ID} = 2$ is the most important feature (see Fig. 6), where $f_{ID} = 2$ corresponds to the logarithm of the ratio of barycentric period and dispersion measure. An important point to note is that AdaBoost, XGBoost and GBC have fewer hyper-parameters than ANN (see Section 6 and Section 7). Hence, this shows that training AdaBoost, XGBoost and GBC is a lot easier than training a Neural Network.

In the case of ANN, Type II errors ('False Negatives') are much more than in the case of AdaBoost, XGBoost, and GBC. This proves that the performance of these three algorithms is superior compared to ANN. One would notice that AdaBoost, XGBoost, and GBC perform almost identically in three-fold cross-validation and in the feature importance. We have an analogy to present which explains this. *Consider a teacher and her student. If the teacher wants the student to learn an abstruse lesson, she can either have as many quizzes as possible, wherein she asks about the same lesson, or have only one quiz with only one question from that lesson.* Here, the strategy of conducting multiple quizzes is how Gradient Boosting Classifier works, and the strategy of conducting only one quiz with one question is how AdaBoost works.

Regarding the feature selection, which was done in Section. 4, we report that feature importances as calculated by AdaBoost and XGBoost are in agreement with each other. However, the mRMR approach, which is model independent, yields feature importances more in agreement with GBC. It is intriguing that Adaboost and XGBoost yield similar feature ranking despite the internal differences they possess.

An ideal classifier must be able to classify both pulsars and RFI correctly. In other words, it should have zero misidentified pulsars and zero misidentified RFI events. Here, we try to compare performances of the four machine learning algorithms by plotting the error rates as a function of threshold probabilities.

Each machine learning technique we have considered here, predicts the probabilities of a candidate belonging to both the classes. We record these probabilities, when we perform 3-Fold CV and then plot the error rate of pulsars and RFI for each threshold probability. These plots can be found in Figures 9, 10, 11, and 12 for ANN, AdaBoost, GBC, and XGB respectively. We have also displayed the zoomed-in versions of the same figures in the right panels for better resolution.

An interesting point to observe is that while AdaBoost and GBC seem to perform in a similar fashion, when we compare their confusion matrices (Figures 3 and 5), AdaBoost proves to perform much better, when we compare the error rates versus threshold probabilities (Figs. 10 and 11). However, since XGBoost, like GBC, is a variant of Gradient Boosting, we find similar error rates versus threshold probabilities (see Fig. 12), when we compare metrics, we find the performances of Adaboost and XGBoost to be very similar (see Table 9).

In any classifier model, there are two degenerate cases where the classifier labels every observation as positive or negative, irrespective of the data. We achieve 100% Recall in the former and 100% Precision in the latter case. The overlap is defined as the area under 'blue' and 'green' lines in plots such as Figure 9, showing error rates as a function of threshold probability. This area takes into account the range of thresholds between these two degenerate cases and also the error rate we are achieving for the thresholds in the average sense. Considering only one of the above two is not sufficient since we can have cases where the range of threshold probabilities spans over the entire [0,1] range, but error rate achieved is monotonic (or slowly varying). This is the case when our model outputs probabilities with a large dynamic range. Another case to consider is when the range of thresholds is very small but the error rates obtained are very high. Such a case arises when a model is outputting probabilities, which are not segregated enough.

In the case of a pristine classifier, we expect to have minimum or no overlapping area under the blue and green lines in Figure 9 and similar plots. Hence, to determine the best of the four machine learning algorithms tested here, we compare this overlapping area amongst them. We have used Simpson's rule to perform numerical integration of the overlapping area and tabulated our findings in Table 11. No other paper has used such a metric before, to the best of our knowledge to rank machine learning models.

As we can see, Adaboost outperforms everyone else and has the least overlapping area of 2.76×10^{-6}. This agrees with the results found by the optical astronomy community (eg. Hoyle et al. 2015), who find that Adaboost is superior compared to other supervised machine learning techniques.

We hope that our results based on publicly available post-processing data from the HTRU-S survey processed with the `PEASOUP` pipeline provide impetus to the pulsar community to consider Adaboost and XGBoost in their pulsar (and other radio transient search) pipelines to aid in the automated separation of RFI and other noise candidates. Conversely, we hope this work spurs other pulsar groups to make their post-processing data public (similar to done for optical large-scale structure photometric and spectroscopic surveys), so that people outside the pulsar community can run and tune their favorite machine learning algorithms on these datasets.

ACKNOWLEDGMENTS

We would like to thank Vincent Morello and Ewan Barr for useful correspondence about the HTRU dataset and the SPINN algorithm and for making their data public. We are also grateful to P.K. Srijith and Ben Hoyle for many enlightening discussions about machine learning and especially to Ben Hoyle for emphasizing to us the advantages of AdaBoost.

REFERENCES

Acquaviva V., 2016, MNRAS, 456, 1618
Baghram S., Afshordi N., Zurek K., 2011, Phys. Rev. D, 84, 043511
Baldi P., Brunak S., Chauvin Y., Andersen C. A. F., Nielsen H., 2000, Bioinformatics, 16, 412

Table 10. Precision, Recall, and False Positive Rate (FPR) of the four machine learning algorithms applied to the HTRU-S dataset. Definitions of these metrics can be found in Table 2. FPR corresponds to the blue curve in Figures 9, 10 and 11. In all the four machine learning algorithms, we have judiciously selected the three thresholds P_T and computed the recall and FPR for these values.

	ANN			Adaboost			GBC			XGBoost		
Precision	0.93	0.94	0.97	0.996	0.999	1.0	0.994	0.995	0.996	0.995	0.9964	0.9969
Recall	1.0	0.99	0.96	1.0	0.74	0.20	1.0	0.998	0.991	1.0	0.996	0.985
P_T	0.5	0.7	0.9	0.505	0.515	0.530	0.84	0.9	0.94	0.85	0.91	0.94
FPR(%)	1.5	1.2	0.5	0.077	0.013	0.0	0.1	0.09	0.08	0.08	0.07	0.06

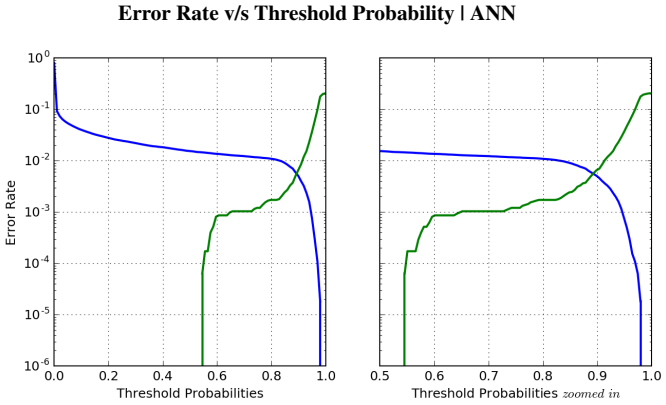

Figure 9. Error rate as a function of threshold probability for ANN. ANN outputs the probabilities for each observation of the test data. For a given threshold probability, we calculate the classification error. We have considered both types of errors, viz. pulsars misclassified as RFI and vice versa. The left panel shows the error rate for both RFI (blue curve on the left) and pulsar (green curve on the right) as a function of the probability over the entire range a threshold can take. The right panel shows the zoomed version of the plot on the left. The area under the blue and green lines should ideally be zero for a pristine classifier.

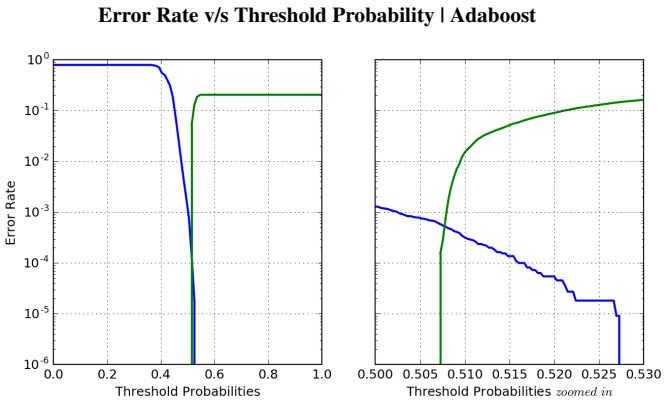

Figure 10. Error rates as a function of threshold probability for Adaboost. More details on the figures in both the panels can be found in Fig. 9. Among all the four tested algorithms, the overlap between the area under the error rate curves is minimum for Adaboost.

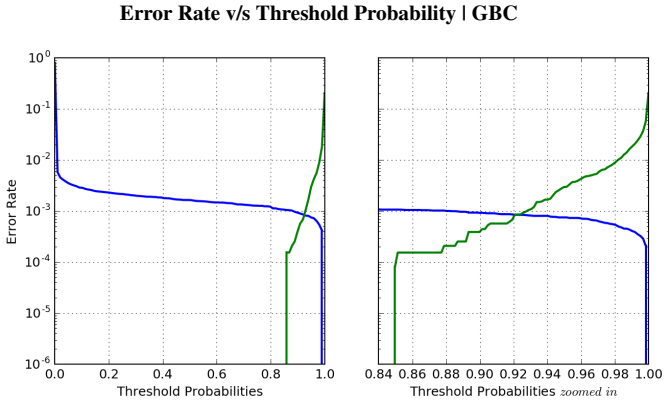

Figure 11. Error rates as a function of threshold probability for GBC. More details on the figures in both the panels can be found in Fig. 9.

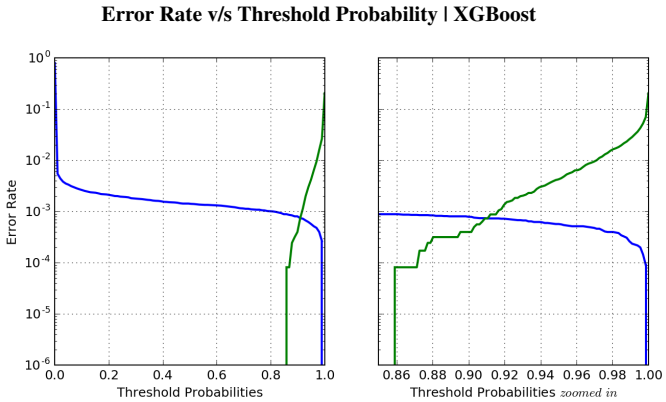

Figure 12. Error rate as a function of threshold probability for XGBoost. More details on the figures in both the panels can be found in Fig. 9.

Table 11. Overlapping area under 'Misplaced RFI's and 'Misplaced Pulsars' curves. We have used Simpson's rule to carry out numerical integration of the overlap and tabulate our findings. A pristine classifier ideally must have zero overlap. In the header column, we define `AU` to be the acronym for "Area Under". See Figures 9, 10, 11, 12, and 12 for visualizing the overlap. *Please note that no other paper has used such a metric before to the best of our knowledge.*

Model	`AU` **Misplaced RFI**	`AU` **Misplaced Pulsar**	`AU` **Overlap** $\times 10^{-6}$
XGBoost	9.64×10^{-5}	0.0011	61.5
Adaboost	9.45×10^{-6}	0.0018	2.76
GBC	0.001	0.0007	73.6
ANN	0.004	0.007	675

Ball N. M., Brunner R. J., 2010, International Journal of Modern Physics D, 19, 1049
Bates S. D., et al., 2012, MNRAS, 427, 1052
Belanich J., Ortiz L. E., 2012, preprint, (arXiv:1212.1108)
Bishop C. M., 2006, Pattern Recognition and Machine Learning (Information Science and Statistics). Springer-Verlag New York, Inc., Secaucus, NJ, USA

Blandford R. D., 1992, Philosophical Transactions of the Royal Society of London Series A, 341, 177
Cameron A. D., Barr E. D., Champion D. J., Kramer M., Zhu W. W., 2017, preprint, (arXiv:1703.05581)
Cavuoti S., 2013, preprint, (arXiv:1304.6615)
Chen T., Guestrin C., 2016, preprint, (arXiv:1603.02754)
Desai S., Kahya E. O., 2016, Modern Physics Letters A, 31, 1650083
Desai S., et al., 2012, ApJ, 757, 83
Detweiler S., 1979, ApJ, 234, 1100
Devine T. R., Goseva-Popstojanova K., McLaughlin M., 2016, MNRAS, 459, 1519
Duchi J., Hazan E., Singer Y., 2011, J. Mach. Learn. Res., 12, 2121
Eatough R. P., Molkenthin N., Kramer M., Noutsos A., Keith M. J., Stappers B. W., Lyne A. G., 2010, MNRAS, 407, 2443
Elorrieta F., et al., 2016, A&A, 595, A82
Faucher-Giguère C.-A., Kaspi V. M., 2006, ApJ, 643, 332
Freund Y., Schapire R. E., 1995, in European conference on computational learning theory. pp 23–37
Friedman J. H., 2001, Ann. Statist., 29, 1189
Hastie T., Tibshirani R., Friedman J., 2001, The Elements of Statistical Learning. Springer
He D., Rish I., Haws D., Teyssedre S., Karaman Z., Parida L., 2013, preprint, (arXiv:1310.1659)
Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, Nature, 217, 709
Hoyle B., Rau M. M., Zitlau R., Seitz S., Weller J., 2015, Monthly Notices of the Royal Astronomical Society, 449, 1275
Johnston S., Karastergiou A., 2017, preprint, (arXiv:1702.03616)
Kaspi V. M., Kramer M., 2016, preprint, (arXiv:1602.07738)
Keith M. J., Eatough R. P., Lyne A. G., Kramer M., Possenti A., Camilo F., Manchester R. N., 2009, MNRAS, 395, 837
Keith M. J., et al., 2010, MNRAS, 409, 619
Kramer M., Stappers B., 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 36
Lee K. J., et al., 2013, MNRAS, 433, 688
Lorimer D. R., Kramer M., 2004, Handbook of Pulsar Astronomy
Macedo F., Rosário Oliveira M., Pacheco A., Valadas R., 2017, preprint, (arXiv:1701.07761)
Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, AJ, 129, 1993
Matthews B. W., 1975, Biochimica et Biophysica Acta (BBA)-Protein Structure, 405, 442
Mayr A., Binder H., Gefeller O., Schmid M., 2014, preprint, (arXiv:1403.1452)
Mirabal N., Charles E., Ferrara E. C., Gonthier P. L., Harding A. K., Sánchez-Conde M. A., Thompson D. J., 2016, ApJ, 825, 69
Möller A., et al., 2016, J. Cosmology Astropart. Phys., 12, 008
Morello V., Barr E. D., Bailes M., Flynn C. M., Keane E. F., van Straten W., 2014, MNRAS, 443, 1651
Natekin A., Knoll A., 2013, Frontiers in Neurorobotics, 7, 21
Ng C., 2013, IAU Symp., 291, 53
Pedregosa F., et al., 2011, Journal of Machine Learning Research, 12, 2825
Peng H., Long F., Ding C., 2005, IEEE Trans. Pattern Anal. Mach. Intell., 27, 1226
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical recipes in C. The art of scientific computing
Ransom S. M., 2013, in van Leeuwen J., ed., IAU Symposium Vol. 291, Neutron Stars and Pulsars: Challenges and Opportunities after 80 years. pp 3–10 (arXiv:1211.3138), doi:10.1017/S1743921312023046
Sadeh I., Abdalla F. B., Lahav O., 2016, PASP, 128, 104502
Sevilla-Noarbe I., Etayo-Sotos P., 2015, Astronomy and Computing, 11, 64
Smith K. M., 2016, preprint, (arXiv:1610.06831)
Srivastava N., Hinton G., Krizhevsky A., Sutskever I., Salakhutdinov R., 2014, J. Mach. Learn. Res., 15, 1929
Tamayo D., et al., 2016, ApJ, 832, L22
Taylor Jr. J. H., 1994, Reviews of Modern Physics, 66, 711
Tramacere A., Paraficz D., Dubath P., Kneib J.-P., Courbin F., 2016, MNRAS, 463, 2939
Wolszczan A., Frail D. A., 1992, Nature, 355, 145
Yang H.-J., Roe B. P., Zhu J., 2005, Nuclear Instruments and Methods in Physics Research A, 555, 370
Yusifov I., Küçük I., 2004, A&A, 422, 545
Zevin M., et al., 2017, Classical and Quantum Gravity, 34, 064003
Zhang H., 2004, in Barr V., Markov Z., eds, Proceedings of the Seventeenth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2004). AAAI Press
Zhu W. W., et al., 2014, ApJ, 781, 117
Zitlau R., Hoyle B., Paech K., Weller J., Rau M. M., Seitz S., 2016, MNRAS, 460, 3152